DOCUMENT RESUME

ED 314 045 IR 052 967

TITLE Linked Systems.
INSTITUTION Association of Research Libraries, Washington, D.C.
REPORT NO ISBN-0-918006-14-7
PUB DATE 6 May 88
NOTE 39p.; Papers presented at the Membership Meeting of the Association of Research Libraries (112th, Oakland, California, May 6, 1988). They will also be included in the proceedings of the meeting.
PUB TYPE Information Analyses (070) -- Reports - Descriptive (141) -- Speeches/Conference Papers (150)

EDRS PRICE MF01/PC02 Plus Postage.
DESCRIPTORS Academic Libraries; *Bibliographic Utilities; *Computer Networks; *Cooperative Programs; Higher Education; Information Networks; Information Systems; Library Administration; Library Automation; *Research Libraries; *Shared Library Resources
IDENTIFIERS *Linked Systems Project; OCLC; *Open Systems Interconnection

ABSTRACT

Three papers are compiled here for research library directors: (1) "Background: Open Systems Interconnection," in which David F. Bishop provides fundamental background information to explain the concept of the emerging technology of linked systems and open systems interconnection--i.e., an agreed upon standard set of conventions or rules that, when implemented will enable computers of different manufacturers, controlled by different types of internal programs, to communicate with each other as equals with minimal prior negotiation; (2) "LSP (Linked Systems Project) and the Library Community," in which Henriette Avram provides an overview of the current status of LSP, i.e., a program involving the Library of Congress, the American Library Association (ALA) Research Libraries Group, OCLC (Online Computer Library Center), and other groups, that is attempting to develop a means of sharing data from one system to another without having to learn each individual system and having a prescribed terminal for each system; and (3) "LSP: Implications for Our Libraries," in which Dorothy Gregor imparts a sense of what linked systems will mean for the future operations of their libraries, e.g., being able to link local online catalog systems with the networks so that library users can move from searching the local catalog to searching OCLC. A list of 11 selected readings is included. (5 tables) (SD)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

ED314045

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

X This document has been reproduced as received from the person or organization originating it.

☐ Minor changes have been made to improve reproduction quality.

• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

Linked Systems

"PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY

C.A. Mandel

TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)."

LR052961

Association of Research Libraries

BEST COPY AVAILABLE

Linked Systems

A program

sponsored by

The Association of Research Libraries

May 6, 1988

Oakland, California

ASSOCIATION OF RESEARCH LIBRARIES

WASHINGTON, D.C.

1988

These papers were presented as a program session at the 112th Membership Meeting of the Association of Research Libraries on May 6, 1988 in Oakland, California. The proceedings of the meeting will be published in 1988 as part of the series, *Minutes of the Meetings*.

The paper used in this publication meets the minimum requirements of *American National Standard for Information Sciences--Permanence of Paper for Printed Library Materials, ANSI Z39.48-1984*.

ASSOCIATION OF RESEARCH LIBRARIES
1527 NEW HAMPSHIRE AVENUE N.W.
WASHINGTON, D.C. 20036
(202) 232-2466

ISBN 0-918006-14-7

TABLE OF CONTENTS

Preface v

Background: Open Systems Interconnection 1
David F. Bishop

LSP and the Library Community 17
Henriette Avram

LSP: Implications for Our Libraries 27
Dorothy Gregor

Selected Readings 33

PREFACE

As technology advances, it is essential that the directors of large research libraries keep abreast of its opportunities and limitations. The program session at the 112th ARL Membership Meeting, held May 6, 1988, in Oakland, California, gave the directors of ARL's member libraries a closer look at one aspect of technology that has had and will continue to have significant impact on research libraries operations: Linked Systems.

The three papers in this publication were all presented at the May 1988 ARL Meeting. The presentations were designed to offer research library directors the background information they need to understand the emerging technology of link systems and related issues, an update on the current Linked Systems Project, and a sense of what linked systems will mean for the future operations of their libraries.

The list of selected readings included in this publication was compiled with the assistance of Beacher Wiggins of the Library of Congress.

The papers presented in *Linked Systems* will also be included as part of ARL's regular series: *Minutes of the Meetings*. The *Minutes* of the 112th Meeting will be published later in 1988.

Nicola Daval
ARL Program Officer
August 1, 1988

BACKGROUND:

OPEN SYSTEMS INTERCONNECTION

DAVID F. BISHOP
Director of Libraries
University of Illinois

We will talk today about linked systems. Because this is a large topic and because there is not a great deal of time available, we have set limited objectives for ourselves. We have tried to focus on the points we believe will most likely be of interest and of value to you as library administrators.

The linked systems project is a complex enterprise with many players involved in its development. It seemed in doing the planning for this program that we had two choices. The first choice was to involve a majority of the players responsible for the development of the Linked Systems Project (LSP) and present eight or nine very brief reports. The result of that approach, I am afraid, would have been most of you going away more confused than enlightened. The second choice was to deal with the more fundamental concepts of linked systems in some depth and in the process ignore--or perhaps a better way to state it would be to fail to acknowledge--some very important contributions that have been made to the development of linked systems.

We have decided to deal with concepts. But before exploring these concepts, I would like to take a minute and describe some of the contributions that have been made to the Linked Systems Project. I hope this will give you a sense of the breadth of the involvement that has occurred.

In looking at contributions to linked systems one would have to begin with the Council on Library Resources (CLR) and Bibliographic Services Development Program (BSDP), which was directed by Lee Jones. Certainly, one of the important contributions of the Council was significant financial support. But equally important was the leadership and vision that existed and that continues to exist within the Council.

Another major contributor is the Library of Congress (LC). The Library of Congress shared in CLR's leadership and vision, and had a great deal to do with initiating the project. And of course, without LC's commitment, the success of the Linked Systems Project would have been highly unlikely.

The Research Libraries Group, Inc. (RLG) has been a major participant. The thing that is most impressive to me about RLG's contribution is that a significant commitment of RLG's resources, both of funding and staff time, were made when those resources were extremely scarce. The linking of systems has clearly been a high priority for those at RLG.

OCLC joined the Linked Systems Project after it had begun and has become a major contributor. If the leaders at OCLC were paranoid, they could easily be convinced that the timing of the development of linked systems was calculated to place OCLC in the worst position possible. OCLC has done a commendable job of finding creative ways of fitting linked systems activities into a massive systems redesign effort.

The National Library of Canada is another participant. A number of important studies have been done by the National Library of Canada, as has some particularly significant work in the area of interlibrary loan and standards for interlibrary loan messaging.

There have been local systems involved in linked systems. For example, GEAC has been an active participant in the project and has successfully established and maintained a link between its system at New York University and RLG. NOTIS is another local system. Jim Aagaard was heavily involved with the initial software development for the Linked Systems Project, and is working on developing a general linking capability for NOTIS users that could be used by other IBM-based systems as well.

And finally there is AVIAC. AVIAC stands for the Automated Vendors Industry Advisory Committee. The major area that AVIAC has focused on is the development of a standard for patron records to allow circulation systems of different vendors to interwork.

You can see that, while we will talk about a limited number of aspects of linked systems, the overall effort has involved a substantial commitment of time and funding by a large number of organizations and individuals.

Before I address my part of the program, which is the open systems interconnection, I would like to take a minute and tell you why the topic

of linked systems is worth two hours of your time. I believe that the impact that linked systems will have on libraries during the next twenty years will be as great or greater than the impact the MARC formats have had during the past twenty years. I believe that the linking of systems will transform the way libraries function and will alter fundamentally the relationship among libraries.

It is as difficult for us to envision the possible uses of the linked systems capabilities today as it was for us to envision the uses of MARC in the late 1960s. In the '60s, we could understand the value of computer-produced catalog cards (they could be customized) and we had vague ideas about what we might be able to do with the machine-readable records. We are in that same position today with linked systems. We can understand the value of a direct connection from our local systems to our utilities rather than having duplicate terminals and engaging in clumsy downloading through printer ports, but we cannot conceive of much more. And, if understanding the uses of linking systems is difficult, understanding the impact that those uses can have on the way libraries function is many times more difficult.

There are a number of parallels between MARC and linked systems. In each instance the development of MARC was, and of linked systems is, ahead of the current technology. With MARC, the kinds of applications that we take for granted today would have been completely impossible in terms of hardware capacity twenty years ago. Also, the structure of the format presented some real programming challenges, particularly for programmers who worked only in high level languages. Today with linked systems we do not have the standard off-the-shelf vendor software we need. We do not have the telecommunications networks we need. There would not be much of a market for off-the-shelf software even if we had that software because the necessary relationships for linking do not exist. And finally, until recently we have not had a complete open systems interconnection standard.

There is an advantage to being ahead of the technology and it is one of the positioning. We saw this with MARC. Because of the early development of the MARC formats, we in libraries do not have to contend with a half dozen competing bibliographic formats. Hopefully, the aggressive way in which the linked systems participants have moved will mean that we will avoid conflicting telecommunications conventions as well.

Now that I have, I hope, convinced you that this topic is worth your time, let me turn to my part of the program, which is the open systems

interconnection, or OSI. OSI is an agreed upon standard set of convention or rules that, when implemented, will enable computers of different manufacturers, controlled by different types of internal programs, to communicate with each other, as equals, with minimal prior negotiation.

There are four key concepts in that definition that I would like to discuss. The first is that of a standard. The specifications for OSI are agreed upon internationally by the International Standards Organization, or ISO. As our library relationships expand, the fact that OSI is not only a standard but an international standard becomes significantly more important.

The second concept is that of different computer manufacturers and software. Computers of the same manufacturer have communicated with one another for a long time. SNA, which is IBM's product, is an example of a single manufacture's telecommunication system. IBM says that SNA stands for Systems Network Architecture. There are others, though, who say it stands for Stop Non-IBM's Applications. This points out an advantage of OSI. OSI does not give large manufacturers an advantage nor does it exclude smaller ones. The result is that OSI frees libraries to select systems and computers without worrying about whether those computers will be able to communicate with computers in other libraries.

Another way in which computers have communicated with one another for some time is by using emulation programs. Emulation programs do the type of thing that OSI does from one particular computer to another. These emulation programs connect computers in a nonstandard way by making one type of computer look like another type of computer. The problem with this approach is obvious. One would have to have as many different emulation programs as there are types of computers with which one wishes to communicate, a sort of telecommunications tower of Babel.

The third concept I would like to discuss is that of equals. OSI does not establish a slave-master or terminal-computer relationship. It establishes a relationship between equals, with each computer regarding the other as a peer. We will see later that this equality requirement complicates things considerably.

The fourth concept is that of minimal prior negotiation. The way I like to think of the capability of OSI is that it is like speaking to someone on a telephone. I can dial anywhere in the world and communicate with minimal prior agreement. Among the prior agreements that I would have to have is that we speak English. The ability to have one computer

communicate with another computer spontaneously is a powerful capability that will be increasingly important in the future.

Finally, I would point out that OSI is not a library standard. Each application group--for example, libraries--will take from the standard certain aspects and create an application implementation. Henriette Avram will be discussing the library implementation later.

Now let us look at OSI in some more detail. At this point I must introduce the difficult part, which is layering. I wish that in presenting OSI I could ignore layering completely, but the minute people begin discussing OSI they talk about things like "the status of the session layer standard" or "the availability of transport layer software," and so for that reason alone we need to spend some time on layering.

As background for layering, I would like to look at structures of library applications programs and then relate those structures to OSI. Chart 1 (see below) shows a multi-application library system with a controlling program

CHART 1

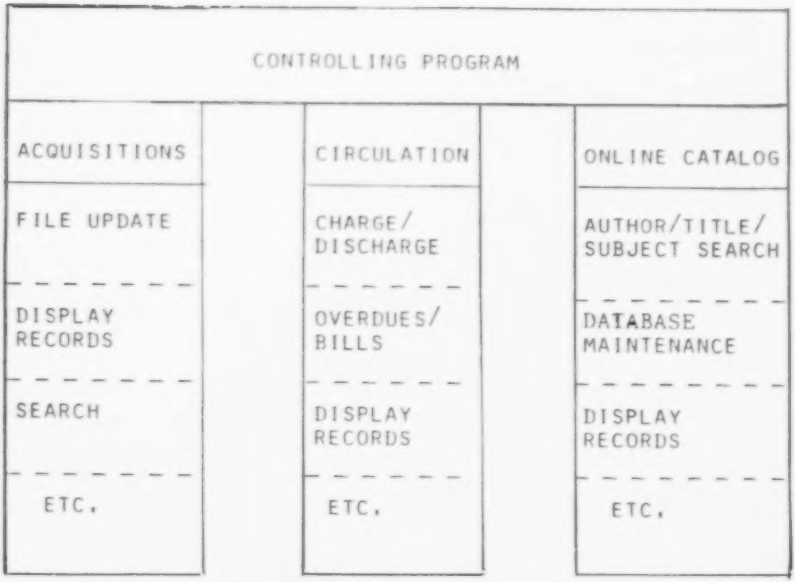

CHART 2

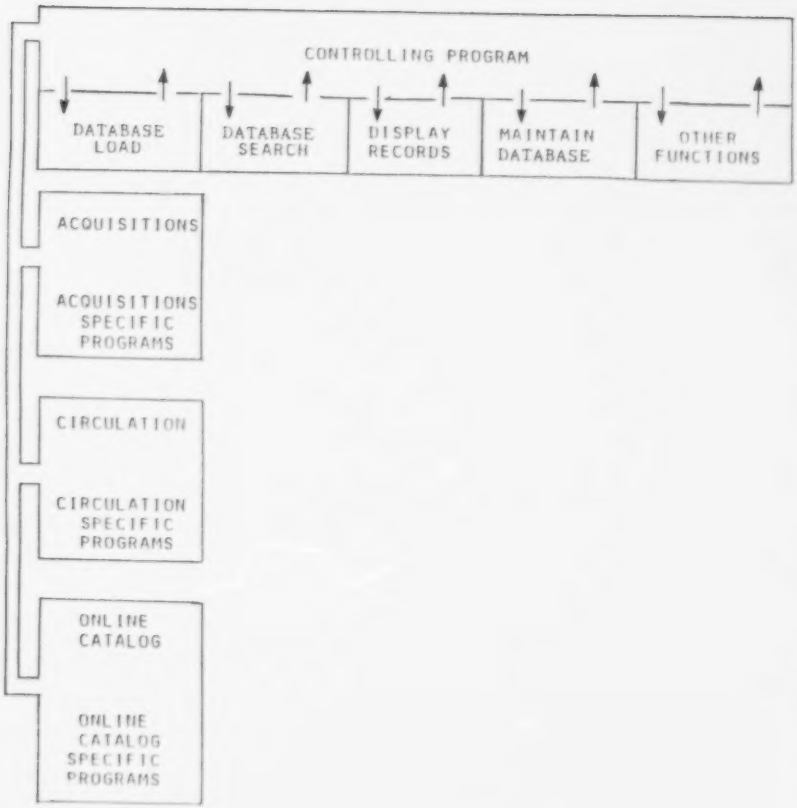
CONTROLLING PROGRAM
DATABASE LOAD
DATABASE SEARCH
DISPLAY RECORDS
MAINTAIN DATABASE
OTHER FUNCTIONS
ACQUISITIONS
ACQUISITIONS SPECIFIC PROGRAMS
CIRCULATION
CIRCULATION SPECIFIC PROGRAMS
ONLINE CATALOG
ONLINE CATALOG SPECIFIC PROGRAMS

that is responsible for the three subsystems: acquisitions, circulation, and the online catalog. In the early days of multi-application library systems, each subsystem or application tended to be self-contained. Programs within the subsystem--such as file update, display record, and search--were written in an individual and often idiosyncratic way, and programs that did identical functions in different subsystems were, in many cases, completely different. For example, in Chart 1 there are three display records programs, one for each subsystem, and each display program probably was completely different. There was one advantage of this approach: a change in the circulation subsystem did not affect the acquisitions or the online catalog subsystem. But, in most cases, the nature of the change required that more than one subsystem be altered, and this resulted in having to change similar but different programs in a number of subsystems.

A better way to organize applications programs is to generalize the software. This approach uses programs that cross applications as shown in Chart 2 (see page 6). Here, the controlling program has a number of generalized routines or programs associated with it. These programs include such things as database load, database search, display records, and maintain database. Connected to the controlling program are the three subsystems: acquisitions, circulation, and the online catalog. These subsystems use the generalized routines through the controlling program to perform functions that are common to all subsystems. This results in less software and software that is easier to change because each routine does a specific, defined task and has a controlled number of entrances and exits. The controlled entrances and exits may be the most important feature of this approach. If one replaces the database search routine, one can merely test the information that enters and leaves the new database search routine and in that way assure oneself that the program is functioning properly.

Now, let us return to OSI. Layering in OSI is similar to the generalized applications program approach in that each layer does a specific task and has controlled entrances and exits. But there is a difference because the OSI layers are hierarchical with each layer connected to another layer rather than having the layers managed by an overall controlling program.

The seven OSI layers, which are shown in Chart 3 (see page 8), are the application layer, the presentation layer, the session layer, the transport layer, and the bottom three layers--the network layer, the link layer, and the physical layer. We will not be concerned with the bottom three layers because in the present LSP implementation, these layers are a standard

CHART 3

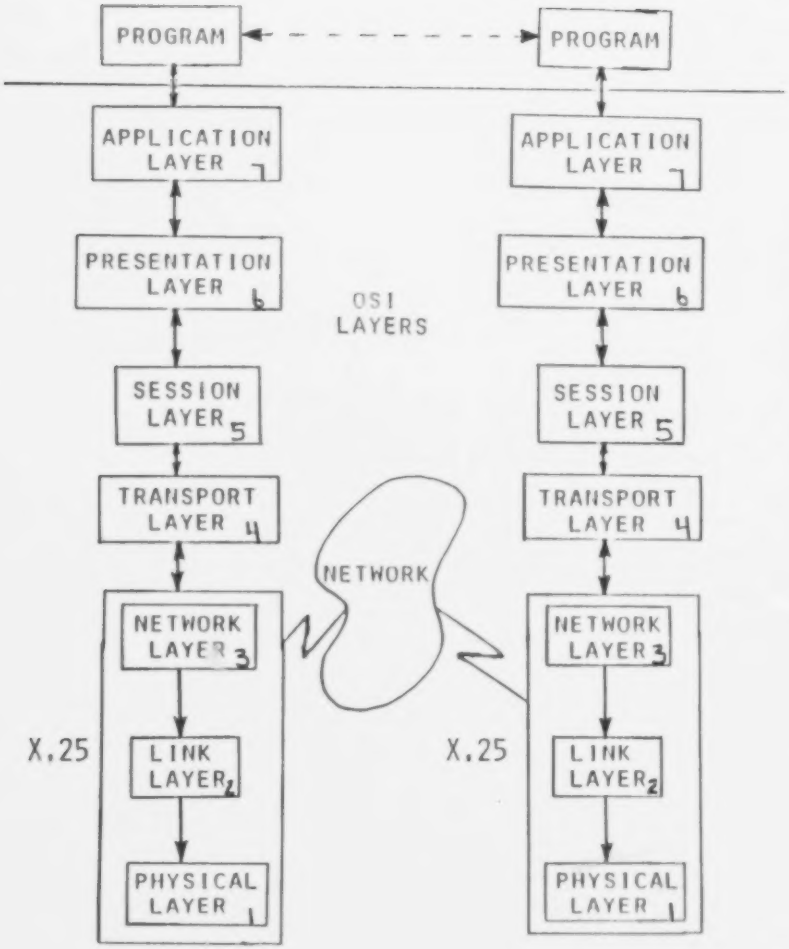
PROGRAM
PROGRAM
APPLICATION LAYER 7
APPLICATION LAYER 7
PRESENTATION LAYER 6
PRESENTATION LAYER 6
OSI LAYERS
SESSION LAYER 5
SESSION LAYER 5
TRANSPORT LAYER 4
TRANSPORT LAYER 4
NETWORK
NETWORK LAYER 3
NETWORK LAYER 3
X.25
LINK LAYER 2
X.25
LINK LAYER 2
PHYSICAL LAYER 1
PHYSICAL LAYER 1

which is called X.25. Examples of X.25 networks are TELENET, TYMNET, UNINET and, maybe someday in the future, LIBNET.

The flow of the information is from layer to layer. The two programs at the top of Chart 3 communicate with each other by sending a message down through the layers until it gets to the X.25 portion. Then, it travels across the physical network to the receiving computer and up through the layers until it reaches the receiving program.

Now we will look at Chart 4 (see page 10), which shows the functions of each layer in more detail. It is unlikely that many readers will want to remember the functions of each of the layers; I have difficulty remembering all of them except when I am working with them. But it is important to understand the type of thing that each layer does.

The application layer, in a library application, contains library-related programs. These programs do such things as taking search requests and translating them into intersystem search formats. The presentation layer does data transformation. It is responsible for translating data from EBCDIC to ASCII, if that is necessary. It could also do data compression, and for applications that require security, it could do encryption. Also the presentation layer handles the representation of graphic information such as videotext.

The session layer keeps different application programs separate, thus allowing a number of programs to share the same communications path. There might be a search program, an interlibrary loan request program, and a record transfer program all being executed at the same time. The session layer keeps those programs separated so that they do not become confused. A good way to think of the session layer is that it connects programs the way the lower layers connect computers.

The transport layer does three things. It handles multiplexing, which I will not talk about, error detection and recovery, and flow control. Error detection is fairly self-explanatory. The transport layer finds transmission errors. It accepts messages with errors that are found by the lower layers and, in some cases, it attempts to recover from errors.

Flow control is an interesting part of the transport layer and is a good example of the problems that arise when two equal or peer computers attempt to communicate. Flow control between a computer and a terminal is simple. The computer controls the flow; we call it response time. With equal or peer computers that is not the case. It would be possible to have

CHART 4

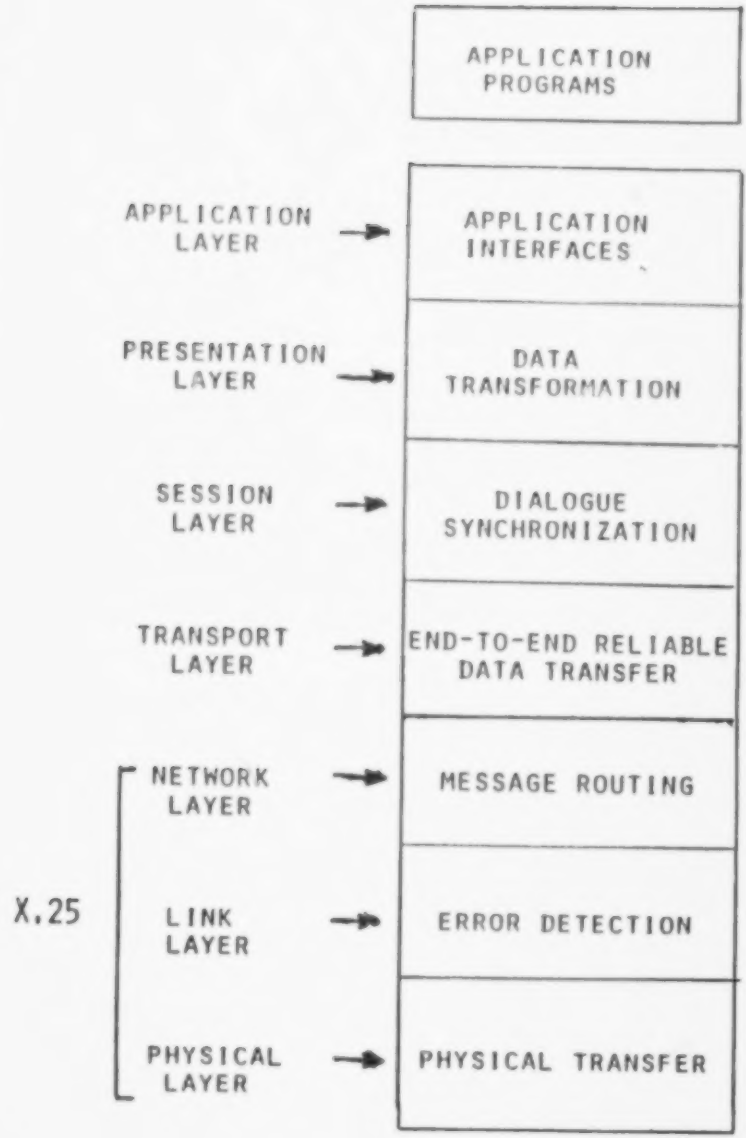
APPLICATION PROGRAMS
APPLICATION LAYER
APPLICATION INTERFACES
PRESENTATION LAYER
DATA TRANSFORMATION
SESSION LAYER
DIALOGUE SYNCHRONIZATION
TRANSPORT LAYER
END-TO-END RELIABLE DATA TRANSFER
NETWORK LAYER
MESSAGE ROUTING
X.25
LINK LAYER
ERROR DETECTION
PHYSICAL LAYER
PHYSICAL TRANSFER

three computers sending data to the same receiving computer as fast as possible all at the same time, which the receiving computer probably could not handle. Flow control regulates the amount of data that can be transmitted in the following manner. The sending computer must ask permission of the receiving computer to send messages. The receiving computer authorizes a given number of messages, after which the sending computer sends those messages and then asks permission to send more. The receiving computer controls the amount of information that is received by controlling the number of messages that it will allow to be sent. It is the transport layer that manages the flow data. The network, link, and physical layers as I said earlier we will not consider.

Now I would like to look at the relationship between the sending and the receiving layers as shown in Chart 5 (see page 10). In Chart 5, there is a user query that is sent to the application layer or entity, and the application entity formulates a search protocol message. That search protocol message is intended for the receiving application entity, and the dashed line between the two application entities shows this relationship. The application information that is to be transmitted to the receiving application entity is the APDU, or application protocol data unit. The APDU is sent down to the sending presentation layer where the presentation entity adds the presentation protocol control information. That message is to be communicated to the receiving presentation entity through a PPDU, or presentation protocol data unit. Next, the message moves down to the session layer, or entity, where the same general process is repeated. When the message reaches the network layer, the X.25 information is added, and it is sent across the physical network and up the layers on the receiving side, where in each case the layer strips off the information that is intended for it and sends the remaining message to the next higher layer, where more data is stripped off until the message finally reaches the receiving program. When the message reaches the receiving program, it looks exactly the same as it did when it left the sending program. Each layer, then, is communicating with its peer or its comparable layer. The process is an add-on process where the higher level message is encapsulated in the lower level message until it is actually transmitted and then the process is reversed. The important thing to remember is that each of these layers performs a fairly simple task in the overall communication process, receiving messages from one layer, and passing them on to the next. But in fact, each layer is communicating with its peer or comparable layer.

So to review, OSI is a set of rules that dictate, in detail, the conventions that are necessary for computers of different manufacturers,

CHART 5

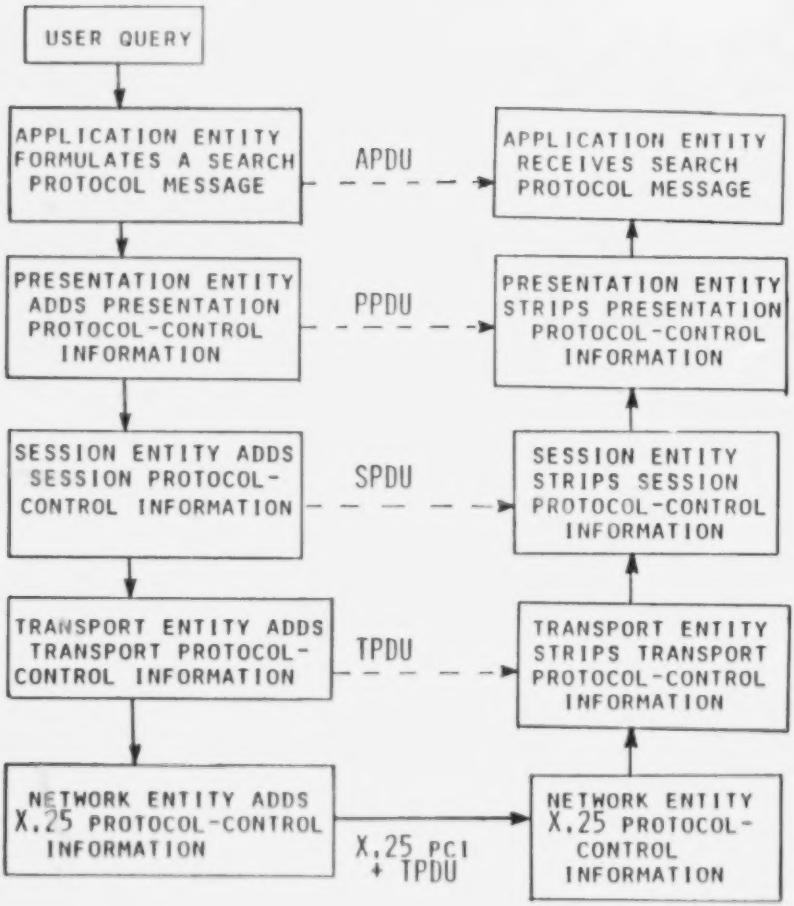
USER QUERY
APPLICATION ENTITY FORMULATES A SEARCH PROTOCOL MESSAGE
APDU
APPLICATION ENTITY RECEIVES SEARCH PROTOCOL MESSAGE
PRESENTATION ENTITY ADDS PRESENTATION PROTOCOL-CONTROL INFORMATION
PPDU
PRESENTATION ENTITY STRIPS PRESENTATION PROTOCOL-CONTROL INFORMATION
SESSION ENTITY ADDS SESSION PROTOCOL-CONTROL INFORMATION
SPDU
SESSION ENTITY STRIPS SESSION PROTOCOL-CONTROL INFORMATION
TRANSPORT ENTITY ADDS TRANSPORT PROTOCOL-CONTROL INFORMATION
TPDU
TRANSPORT ENTITY STRIPS TRANSPORT PROTOCOL-CONTROL INFORMATION
NETWORK ENTITY ADDS X.25 PROTOCOL-CONTROL INFORMATION
X.25 PCI + TPDU
NETWORK ENTITY X.25 PROTOCOL-CONTROL INFORMATION

using different controlling programs, to communicate with each other without or with minimal prior agreement. These protocols are described in layers so that each element or layer is simple enough to be explained easily and so that it can be tested against a standard easily. And, the rules that make up OSI must be implemented in order to actually achieve communications.

I would like to finish my part of the program by addressing some OSI issues. One can think of this section as an OSI survival kit for library directors in dealing with their computer centers, local systems vendors, and network people at the state and regional levels.

The first OSI issue to address is interworking, and it is important to understand what is meant by interworking. There is a parallel between OSI and AACR2. If two libraries had implemented AACR2, and if in each case where AACR2 provides for options the two libraries selected different options, their cataloging would not interwork or in the case of catalog cards would not interfile very well. OSI has options just as AACR2 does. Therefore, interworking, or the ability to actually communicate, requires agreement about which options to chose or a program's ability to accommodate multiple options automatically.

Interworking is clearly necessary if library systems are to communicate with each other. There must be interworking from one vendor's software as a sender to another vendor's software as a receiver. There must also be interworking of different layers by different vendors. For example, one vendor's transport layer software must interwork with another vendor's session layer software. It is dangerous to state the status of anything on a subject that is changing as rapidly as this is, but I know of no example of interworking of off-the-shelf standard OSI software from different vendors. Clearly, interworking must take place before libraries can make much practical use of OSI software. As I said earlier, we are ahead of the technology.

The second OSI issue is efficiency. There are claims that the OSI protocols are inefficient. One author actually calculated the delays that would occur as messages move from one layer to the next and suggested that the protocols would be too slow for many applications. There are two points that should be made with regard to efficiency. First, a protocol is neither efficient, nor inefficient, nor fast, nor slow--only implementations or programs are. The early X.25 programs were inefficient and slow; and through reprogramming they improved. The upper layer implementations of OSI should become more efficient as well.

The second point about efficiency is that layers are not necessary for a particular implementation. Layers are a means of describing and understanding the OSI standard. Ray Denenberg from the Library of Congress addresses this issue in an article in *Library Hi Tech*.[1] He says "the layering technique is used to allow individual groups of experts who are developing protocols to concentrate on well defined subsets of the range of communication functions. However, OSI does not prescribe that the internal implementation of a system reflect layering." He goes on to say "some of the criticism leveled at OSI reflects a misunderstanding of this point. The belief that layering is elegant but too expensive refers to a mistaken interpretation of the OSI model that individual layer entities must be constructed and that the resulting system is overburdened by all of the interactions between layer entities. In fact these distinctions and interactions between layers are only conceptual." So, while efficiency will be a matter of concern for the next little while, there is no reason to believe that efficiency will be a long term problem for libraries or for the overwhelming majority of other computer applications.

The final issue I would like to discuss is TCP/IP. TCP/IP are five letters with which anyone involved with OSI becomes familiar very quickly. TCP/IP stands for Transport Control Protocol/Internet Protocol. TCP/IP does the same type of thing that an implementation of OSI would do. The reason that you should know about TCP/IP is that it is used by many of today's supercomputer networks. National Science Foundation sponsored programs and universities with heavy NSF commitments are where you are most likely to find TCP/IP networks.

There are good things about TCP/IP. First and foremost, software for TCP/IP exists. It is possible to go out and buy off-the-shelf TCP/IP software, which is not the case with off-the-shelf OSI products. Also TCP/IP is a mature protocol and so TCP/IP programs are fairly efficient.

In spite of its availability, there are two major problems with TCP/IP. The first problem is that it is not a standard. There is no guarantee that software from different vendors will interwork and there is no arbitration mechanism when software from different vendors does not interwork. A second problem is that government agencies including the National Science Foundation and the Defense Department have been directed to migrate from TCP/IP to OSI in the near future.

---

[1] Denenberg, Ray. "Open Systems Interconnection." *Library Hi Tech 3*, no. 1, Consecutive Issue 9 (1985): 15-26.

There will most certainly be a general migration from TCP/IP to OSI. One strategy that has been proposed is to take an OSI message and encapsulate it in a TCP/IP message, much way that OSI does in passing messages from layer to layer. Then, over time, the TCP/IP portion of the message would be eliminated, leaving a standard OSI network message.

The migration from TCP/IP to OSI will likely be a slow one largely because there are so many TCP/IP networks in place. It is important that the library community, which has been a leader in the development of OSI, work closely with TCP/IP computer people to develop migration strategies. We in libraries may find ourselves using TCP/IP networks in conjunction with OSI for some time. But in the long run there is little doubt that OSI will prevail and the fact that we in libraries made the commitment to OSI when we did should make our transition relatively painless.

A development that should have a significant impact on the availability of OSI software occurred recently and was reported in the May 5, 1988 issue of *USA Today*.[2] In an article, entitled "High-Tech World Sets Product Standards," an event of considerable significance was described. The article began

> "Mark your calendar: Wednesday, May 4, 1988, was the day the Information Age left its Dark Ages behind. In New York, 12 USA computer executives stood shoulder-to-shoulder and told the world how all their products--representing 80% of industry sales--will exchange information in the months and years ahead."

The article went on to describe how telephone standards were developed in the late 19th century to allow international calling. Then the article said, "Wednesday, the computer industry followed suit. IBM, Digital Equipment, Unisys, Hewlett-Packard and eight others said their computers will conform to a standard called Open Systems Interconnection, development by the International Standards Organization." The article concluded by describing some of the events that have encouraged the shift to OSI. First, it stated that "starting in 1990, all computers purchased by Uncle Sam must conform to OSI." Secondly, William Gates, chairman of software giant Microsoft Corp., said "software products matching the OSI standard will hit the USA market this year." Thirdly, Telenet Vice-President, Stuart Mathison said Telenet's network of OSI-based computer systems, now serving 300,000 users

---

2 Hillkirk, John. "High-Tech World Sets Product Standards." *USA Today* (May 5, 1988): 1B.

throughout the world, is growing "into a truly global messaging network." And finally, "customers will demand OSI products to make computers more reliable, and eliminate the money and time spent mixing and matching computers. By installing OSI computers: General Motors Corp. has cut factory downtime in Pontiac, MI, by 50%; and Boeing Co. expects to shave its computer networking costs by 30%."

This agreement by computer manufacturers is just one more step toward the eventual realization of national and international OSI telecommunication networks.

I have tried to explain the basics of the open systems interconnection, looking particularly at those things that would be of interest to library directors. I believe that the linking of systems will have a dramatic impact on the way libraries function in the years to come. The emergence of off-the-shelf software will be some time away, as will the development of general OSI networks that are available to libraries. We in libraries need to be aware of continuing OSI developments and we need wherever possible to influence those developments. There are, as I indicated at the beginning, many groups and individuals in the library and information science profession who are monitoring the progress of the implementation of OSI and are playing a role in determining how that implementation will evolve. There is little doubt that we will all benefit from their efforts.

LSP AND THE LIBRARY COMMUNITY:

PRESENT STATUS

HENRIETTE D. AVRAM
Assistant Librarian for Processing Services
Library of Congress

As a community, libraries have been involved in networking activities for some years now. What may not be so well-known to you as library managers are the ongoing efforts to lay the framework for a library network supported by technology that has been in existence for some time now. I will describe the project though which the library community, using the Open Systems Interconnection (OSI) which David has described, has marshalled that technology to make certain that the library network grows into a viable and enduring structure that will serve libraries well in the coming years. I am referring to the Linked Systems Project, or LSP. I will discuss what LSP was intended to do and what we are doing now. You as library administrators have a critical role to play in LSP's success as a networking tool for librarians and should therefore be knowledgeable about what is going on.

Networking among libraries has been sustained by a well-established tradition of sharing resources. But beyond the altruism implicit in this long held tradition, there are the escalating costs associated with operating libraries of which all of you are certainly well aware. The high cost of purchasing materials needed by users and the equally high cost of bringing those materials under bibliographic control so that they can be accessed by users have become prime factors driving the library community to seek expedient means to share data through the building of networks.

The availability of cataloging data in machine-readable form facilitated the computer-based systems we know today as bibliographic utilities and most, if not all, of you are members of these utilities. By the beginning of the 1970s, there had emerged four major databases--the three major utilities (OCLC, the Research Libraries Group's RLIN, WLN), and the Library of Congress--with no way, except for LC data, to share the large repositories of data residing on each utility. What was needed was a way to share data from one system to another without the necessity of learning each

individual system and having the prescribed terminal for each system. The development of LSP was aimed at achieving this end.

The Library of Congress, RLG, and WLN (OCLC was to join later and WLN was to become less actively involved) came together in 1980 to continue the development of a program begun in the mid-1970s to meet the challenge of constructing a nationwide system composed of large independent systems, each with its own tailor-made system architecture, that could operate with the efficiencies of a single network, and thereby increase national resource sharing.

Basic to the success of LSP was the selection of a suite of protocols to define the building of the computer-to-computer links. In the development of LSP, as David has indicated, it was decided to embrace the International Organization for Standardization's Open Systems Interconnection (OSI) Reference Model, and to base the LSP communications facility, the Standard Network Interconnection (SNI) on the model.

In addition to the communications component, the other LSP component is the applications component. The two general applications necessary to effect data exchange are Record Transfer and Information Retrieval. Record Transfer enables records of any type and any number to be transferred between systems. It has two aspects: contribution and distribution. Information Retrieval, by supporting intersystem search, permits users of one system (e.g., OCLC) to access a remote system (e.g., RLIN) and to view data found in the remote system on their own system, using the search commands of their own system. This means that the user sees little or no difference between local searches and the resulting displays and LSP searches and displays. The search and response are controlled by the system initiating the search and appear to the user to be quite straightforward, since the search instructions are the ones normally employed by the user. The complex task of translating a search query from an origin system into an intersite syntax and back into the search query of the target system is the responsibility of the computer and is transparent to the user. This should be clearly understood and is one of the most significant aspects of LSP.

Now that I have set the stage by providing some background, it might be useful to enumerate the project priorities that have been set by the present LSP partners to realize their goal of creating a logical network through the interconnection of systems. They are as follows:

1. **Complete the Authorities Implementation.** The Authorities Implementation was selected as the first priority for LSP because the most expensive part of cataloging is creating headings and related cross references for the building of consistent catalogs.

2. **Develop the means to transfer bibliographic records** to facilitate the implementation of a coordinated cataloging program.

3. **Synchronize the local system linkages to the systems of the LSP partners,** i.e., develop local system linkages to OCLC, RLIN, or LC that are identical to the ones linking OCLC, RLIN, and LC.

4. **Develop the means for intersystem searching to support shared cataloging as we know it today,** i.e., the user searches the database of his utility which in turn could search the database of a remote utility to see if an item has been cataloged, and if so, the record will be transferred from the remote system to the user's utility.

5. **Develop the means to support interlibrary loan transactions.**

Priority one, the Authorities Implementation, is basically completed now and a program is in place to take full advantage of its benefits. This program is the National Coordinated Cataloging Operations, or NACO as it is more commonly called, and is housed at the Library of Congress. The NACO operation, which started in 1977, involves the cooperative building of a database of authority records created by 43 libraries participating in the program administered by LC. The authority records they create are submitted to LC in a variety of ways: on worksheets mailed to the Library; via terminals connected to the LC database; and now computer-to-computer via LSP. Under the LSP configuration, staff at other libraries input their records to their utility and the records are then transmitted via the LSP link to the authority file on LC's database. This is referred to as "contribution." LC sends these records, along with its own authority records, over the link to the utilities where all members of the utilities can use them. This is referred to as "distribution." By utilizing the LSP applications Record Transfer (for contribution and distribution) and Information Retrieval (for intersite searching), records are contributed in a timely fashion (within 24 hours) and all members of OCLC and RLIN are

searching against the most up-to-date national authority file; thus the costly duplicate creation of records is significantly reduced.

Distribution of approximately 2,500 authority records daily from LC to the RLIN database has been in operation for over two years now (approximately one million records). OCLC has been receiving the same authority records in this mode for more than a year. Neither utility any longer loads tapes of LC authority records into its database. In June of last year, RLG began sending authority records, contributed by their members via the link to LC to be added to the authority file. Yale and Princeton Universities, and the University of Michigan are the RLG member libraries contributing in this mode. OCLC began in February of this year; Indiana University is the first institution to participate and will be followed by the University of Illinois at Urbana. It is our intention to have all NACO libraries that are members of OCLC and RLIN contribute authority records via LSP.

In the development of the application protocols, care was exercised to tailor them to be generally applicable and not peculiar to a particular application. With some modification and augmentation, then, it is possible to expand the LSP applications to other purposes. Work is already underway to make Record Transfer and Information Retrieval support the exchange of bibliographic records (priority two). This exchange will come about as part of another cooperative venture involving LC and eight selected research libraries. The eight libraries are Yale University, Harvard University, Indiana University, the University of Illinois at Urbana, the Univesity of Michigan, the University of Chicago, the University of California at Berkeley, and the University of Texas. Designated the National Coordinated Cataloging Program (NCCP), its purpose is to contribute standard records to the national database of cataloging records. NCCP will operate similarly to NACO, taking advantage of many procedures currently in place for NACO. In both projects, the participating libraries have agreed to follow LC cataloging practices in creating these records. Participants in NCCP must be NACO members and must create authority records related to the bibliographic records they submit. They also will provide subject headings. LC has been giving personnel from the participating libraries special training for this effort.

Like the authority records created through NACO, these NCCP records will be contributed to the datebase residing at LC, where they will in turn be distributed over the LSP link to the utilities for access by their members. The two-year pilot to test this undertaking began last month operating on a combination of LSP links (for some authority records) and

terminal-to-LC database (for all bibliographic and the remaining authority records). The first NCCP-created records have been received at LC and added to the database. NCCP appears to be off and running!

It is probably worth my taking a few minutes here to explain why it was decided to proceed with NCCP even though LSP for bibliographic records was not available. A lot of momentum had been generated among the participants; significant agreement had been reached among the originators of the program--e.g., members of RLAC (the Research Libraries Advisory Committee to the OCLC), RLG, and LC--as to how the program should operate; and the Council on Library Resources had granted funding to help defray the costs. It was also felt that much could be learned, even in this less than perfect environment, that could be carried over to the full LSP mode. So the planners gritted their teeth and forged ahead, but not without considerable impediments.

The importance of LSP is underscored by our experiences thus far with the NCCP pilot. Based on only one month's duration, the pilot has already demonstrated to all the advantages of LSP and what a boon it would be to the smooth operation of NCCP. LC and the participating libraries have encountered considerable problems in installing LC terminals and printers at the local sites and getting the telecommunications lines connected. Most of these obstacles would be eliminated if LSP links were fully operational, because participants could then use their own system terminals and familiar search commands for searching, input, and update. LC's commitment to LSP has been greatly reinforced by these experiences.

If the pilot proves successful, the switch will be made to the full LSP environment as soon as it is available. To help determine the success of the pilot, CLR has formed a panel of experts in the field of library evaluation and statistical analysis to evaluate the effectiveness of the pilot. The evaluation will be based on input from LC and the eight participating libraries related to their experiences during the pilot. The results of this ambitious evaluation should have implications for all types of future cooperative programs.

The goal of these cooperative projects is to provide national-level records for items acquired by the Nation's libraries. The Library of Congress cannot possibly catalog everything that is required, but it can work with other libraries to increase the aggregate number of high quality bibliographic records available to libraries. Areas of cataloging responsibility have been agreed upon, although it is not mandatory that the responsible library input everything it receives within its area of

responsibility. This division of responsibility is also based on the available cataloging expertise, that is, staff competencies at particular institutions. Rigid allocations of responsibility have tended to cause problems in past cooperative efforts. For the pilot, leeway for overlap is built in, i.e., if an institution receives an item outside its normal scope and needs a cataloging record, and on searching the database finds that it has not been cataloged by the responsible library, the library in need of the record may input it. The results of the pilot are likely to cause the cataloging assignments to change.

LC enters into cooperative arrangements in the interests of efficiency and economy. When resources can be conserved through the selective use of the work already being done elsewhere, a savings results. Further, our remaining resources can be directed to other areas. These considerations have been fundamental to LC's active participation in LSP and have been central to its involvement from LSP's incubation stage back in the mid-1970s.

Agreement has been reached among the LSP partners that the LSP protocols will be used to effect the linkage of local systems to the bibliographic utilities or to each other (priority three). Local systems have been developed among libraries in response to a need to execute more efficiently tasks performed at the local level, such as circulation control or serials check-in. LSP can make possible the uploading of cataloging records (from a local system to a utility) and the downloading of such records (from a utility to a local system). It is through LSP that these potential "islands" of communication can be bridged, ensuring that resources, and in particular cataloging data, are shared nationwide. New York University's use of a GEAC local system applying the LSP protocols to connect with RLIN is representative of how LSP can be channeled in this direction. Records will be transferred from RLIN to NYU using Record Transfer.

Much work remains to be done to achieve priorities four and five. For priority four, shared cataloging, there is agreement that the member networks (OCLC and RLG) will make intersystem searching available to their member institutions following some agreed upon plan. The National Library of Canada has been working on an interlibrary loan protocol for several years now that will be compatible with LSP requirements and that protocol will most likely be the basis for any LSP interlibrary loan implementation (priority five) in the United States.

As with any undertaking of such magnitude and complexity, it was necessary to have a mechanism for coordinating the activities of the

various LSP players. To this end, several committees were established. The LSP Policy Committee is the overarching forum and as such provides general policy and direction for the linking process. Additionally, it is concerned with costs, accounting, and related administrative matters. It further gives guidance to and receives reports and recommendations from the following two committees.

The LSP Technical Committee was formed with the primary charge of ensuring that the LSP communications protocols conform to the OSI standards and are compatible with OSI implementations. The committee also coordinates any needed changes to implementations, shares expertise among LSP implementors concerning those implementations, and promotes and encourages the use of OSI for library networking. David Bishop has chaired this committee since its inception.

The LSP Application Committee plans and coordinates the implementation of applications, develops and maintains the application layer information exchanged between systems, and specifies application requirements. Currently the committee is working on the analysis and specification requirements for implementing computer-to-computer exchange of bibliographic records in book format. LC has been given administrative responsibility, with me as chair, for this segment of activity. Taking an accelerated approach of intensive and frequent meetings at LC, and with the full and active cooperation and participation of the representatives from RLG, OCLC, GEAC, NOTIS, and the Triangle Research Libraries Network, the committee has made tremendous inroads. The requirements documents are nearly finished and the LSP partners are setting timetables for implementing the bibliographic component. Mid-1989 to mid-1990 is targeted as the likely range of implementation dates for LC, OCLC, and RLIN. When this phase of its work is completed, the committee will tackle requirements for exchange of non-book records and further resolution of local system-to-utility processing problems. The first meeting of the committee to address local systems to utility processing is scheduled for June 1988.

David mentioned the differences in approach to building networks and the network infrastructures that are emerging to link university campuses running on different protocols. As more is known about the various other networks such as National Science Foundation Network (NSFNet) and the New York State Education and Research Network (NYSERNet), it becomes increasingly clear that there is a major difference between what their proponents are establishing and what libraries are attempting. They are planning a pilot project in which a scholar will be able to sit at a

workstation and have searching access to data available on the network, including libraries. I do not minimize this effort but ours is a production-oriented network. By that I mean we are searching and creating records online in an interactive mode. We are inputting data to be exchanged, built upon, and otherwise manipulated for diverse purposes. The planners of the education networks are busily wiring the universities of this country, connecting powerful computers to form a national network. But to date they have few applications for the "supernetwork." Libraries are part of these universities; the opportunity must not be missed to work together. My contacts with key personnel associated with these groups has lead me to believe that we can work together--they are eager to have the library community join with them in building the nation's network.

Just last month a meeting was called by invitation. Its principal aim was to begin to investigate how to address the issues involved in establishing a national network. Large organizations are involved in these plans, e.g., IBM, AT&T, New York Telephone, the National Science Foundation (NSF), and EDUCOM. And yet it was very obvious they have just begun to face problems we have struggled with so long--e.g., standards, governance, copyright, etc. This difference notwithstanding, however, with some careful planning and cooperation with the education and information communities, there is no reason that the national networks being built by these diverse constituencies cannot operate together for the mutual benefit of all of the nation's information seekers. For this to happen, the various communities must become keenly aware of what others are doing, what protocols govern their network structures, and seek ways to tie their networks. LC, OCLC, and RLIN are all working towards this end.

We in the library contingent are certainly becoming attuned to prospects and implications of these networks and the different protocols running them. Moreover, we are mindful that we must cooperate in making sure that we can communicate. For instance, we at LC have worked closely with representatives from EDUCOM to try to ensure that library networking and academic networking do not continue to move in different directions. LC participates on the EDUCOM Networking and Telecommunications Task Force Network Resource Committee. We have also formed a local group that meets periodically and includes LC personnel and staff from the University of Maryland, NSF, and the Department of Education. We have dubbed ourselves the Network Resource Group. Our most recent meeting was last month and our next is scheduled for July. We have had detailed discussions with key individuals in the academic networking community on the subject of aligning our architectures for interoperability.

The LSP Policy Committee has recognized the need for cooperation as well by devoting parts of its agenda to these developments. In the committee's view, it is important to illustrate how library networks foster scholarship across all disciplines. It will further stress that it is imperative that the national network be designed to accommodate the transfer of bibliographic records and other library support data in order to be totally serviceable to the entire information community.

LSP: IMPLICATIONS FOR OUR LIBRARIES

DOROTHY GREGOR
University Librarian
University of California, San Diego

David has talked about the OSI side of linking systems and Henriette has provided an overview of the Linked Systems Project. My job is to suggest some of the future directions for the linked systems capability in our libraries by looking briefly at three areas: extending our current use of the link, looking at what the Canadians are doing in protocol development and, finally, to suggest some possibilities for our local systems development.

As Henriette has described, the current use of the linked systems protocols is in support of bibliographic record production facilities among the Library of Congress and the largest of the bibliographic networks. We have begun to use the link for building cooperatively our national authority file through NACO, and we are looking at its use in support of the National Coordinated Cataloging Program, which will focus its cataloging efforts on monographs in the MARC BOOKS format. What is next?

Those of us who have been involved in the frustrations of trying to build the CONSER file with participation from libraries who are not OCLC users will no doubt think immediately of CONSER as an extension of the production facility use of the link. CONSER is already one of the best things on the bibliographic market and to be able to add those non-OCLC libraries which either are trying, have tried, or would like to try to add records to CONSER would add significantly to the national serials database.

We heard earlier in this meeting that we can expect increased funding from the National Endowment for the Humanities (NEH) for preservation microfilming as part of the preservation effort being organized by the National Commission on Preservation and Access. As we gear up to increase the production of preservation masters to save our collections, we are also thinking about the requirements for recording the salvation of another text. OCLC and RLIN have agreed to exchange records for preservation masters, but tape exchange never seems to work smoothly or promptly--and sometimes it does not at all. As microform master

production increases, there is an increasing need to record that production in a more timely fashion in order to avoid duplicating the expense of filming. There is not yet agreement on the record requirements for preservation masters--that is a topic for ARL's Committee on Bibliographic Control--but there is agreement that we need to provide records for our preservation masters as rapidly as possible and to build a national preservation database. As you know, NEH has funded the conversion of the National Register of Microform Masters and RLG has been adding to its database of master microform records at a good rate. I believe we should extend the use of the link systems capability to records for preservation master microfilms in order to build a national preservation database with the Library of Congress as the manager of the file. This use of the link would move our preservation efforts forward and help convince potential sources of funds that we are serious about sharing the results of our efforts.

Another category of records that lends itself to the production facility use of linked systems is the cataloging of Chinese, Japanese, and Korean titles--the CJK records. And to those we might add any of the more exotic languages in which libraries have particular needs to share expertise because there is not enough to go around. CJK, however, is already an organized cataloging process on systems at LC, RLG, and OCLC and, again, the timely sharing of those expensive to create records would assist in making materials in those languages more readily available to our clientele.

You may well have other categories of materials that would be candidates for early use of linked systems. It would be useful to have some conversation about our greatest needs, so we get the most out of this early, limited capability.

In the U.S., largely under the leadership of the Council on Library Resources and the Library of Congress, our focus has been to link part of the cataloging operations of four of our large bibliographic institutions. The Canadians have taken a different approach. The Bibliographic and Communications Network Pilot Project of the National Library of Canada (NLC) is focused on achieving an open, decentralized, voluntary library and information network, operating nationwide within the OSI environment. The National Library's goal is to foster wide acceptance of OSI within our bibliographic sector by making application protocols and sample programs available to any organization wishing to use them--including vendors of library systems.

Four applications have received most of NLC's attention: file transfer, information retrieval, interlibrary loan, and acquisition. The NLC recognizes that the work it is supporting on applications may not end up the OSI standard, but they are tracking and influencing the eventual standard and, in this way, advancing the standards needed to build their open system.

The interlibrary loan application is the most developed. It has been tested and implemented at the National Library of Canada and is considered a standard within Canada. NLC has sent it forward for consideration as an international standard. NLC has also turned over the ILL application to seven vendors and provided seed money to develop the software for the standard in a variety of hardware environments, including GEAC, UTLAS, and PC-based systems. NLC has also developed a conformance tester for the vendors to check out their conformance to the standard. The application should be ready for widespread use in Canada by next year.

The file transfer application is in use at three Canadian libraries, but is not undergoing further development because of NISO's acceptance of the File Transfer Access and Management (FTAM) application protocol developed in the U.S.

Things are also quiet on the acquisition application front because, while the NLC has had some success in getting Canadian libraries organized, the publishing sector seems less amenable to organizing efforts.

For the information retrieval application protocol NLC is looking at the work of NISO's Standards Committee D, developed in the U.S. as part of the Linked Systems Project.

Canada is also working on a "directories" protocol. The directories protocol work is aimed at meeting the need to know what is where in a world of decentralized systems and distributed databases. This is very complex stuff and, again, part of the National Library of Canada's motivation for their work is to track and influence emerging protocols--to help set the agenda for OSI protocol development and to engage the vendors to make the protocols available once they are developed. As you can see, the approach is quite different from that of the U.S.

Now let's look at what we can do with our local systems. Ultimately, the linked systems protocols are one of the tools we will use to make it seem as though there are not so many bibliographic and other systems in the library world. For example, we can think about connecting local

systems to our bibliographic networks. Some of this has already been done in the same kind of record transfer application used for the cataloging projects. A GEAC system at NYU is downloading records from RLIN. The NOTIS system at Northwestern is working to develop the link to upload records from NOTIS to RLIN and download records from RLIN to NOTIS. Again, we're looking at moving bibliographic records around to support cataloging operations.

We can also think about the possibilities of linking local online catalogs with the networks for public service searching. At the moment, not many vendors are working to make the software available on their online catalog systems, although most indicate that the software is "under development" or will be undertaken when standards become available. Vendors respond to market demand, and libraries are not demanding the capability. But if, for example, OCLC offered the linked system capability and all local libraries had to do was buy some software in order to make it possible for their online catalog users to move from searching the local catalog to searching the OCLC (or any other) system, then libraries would probably start demanding the software from their OPAC vendors. Or, think about the possibility of one of the database vendors like BRS being available for searching over the link from your OPAC terminals. Multiple systems could be available from a single terminal. And then, of course, we would need to think about how we would organize and pay for such service.

When enough libraries are ready for that additional search capability, the message will get to the OPAC vendors. The large vendors will no doubt be the first to develop linked system software and I understand that it's a bit easier to do in a UNIX environment, but library demand is required if we want this kind of capability in our OPACs. Vendors need to see that link software is a viable product.

Within the University of California, the makers of MELVYL in the UC systemwide Division of Library Automation are working with a vendor, Innovative Interfaces, to develop a two-phased link between the MELVYL UC Union Catalog and INNOPAC, Innovative Interfaces's online public access catalog. Initially, the link will allow an INNOPAC user to issue a command and bring up the MELVYL system for searching. For example, a user on the San Diego campus, not finding the item in our catalog, would switch over to MELVYL and have all the resources of UC available for searching. And, a MELVYL user could switch over to INNOPAC. In other words, one terminal can access either system, but users will need to know the search software of each system. In this first phase, the OSI

application protocol will be layered over TCP/IP to connect the MELVYL and INNOPAC machines.

In the second phase of the project, we will be working on transmitting searches between the two systems, so that a user does not need to know the searching commands for both systems, but can use the system with which he or she is familiar. In these information retrieval applications, users connect to the remote system through their own system, so they do not need to learn yet another set of system commands. We must, however, recognize that some system capabilities may be lost in translation until and unless we develop some standards for query languages.

We can continue to speculate about future uses of the link--for example, to deliver documents as well as data--but widespread use among disparate systems is not going to happen very fast. A few years ago all wide area networks used different protocols and could not communicate with one another. Today most networks are migrating toward DARPA's TCP/IP with a commitment to use the OSI protocols when they become standardized and available. This is the direction taken by the University of California and Innovative Interfaces, Inc., in their cooperative project. The Corporate on Open Systems has recently announced a commitment to OSI, but the realization of that commitment will take several, if not, many years. Still, the commitment has been made and we can be proud and very relieved that major vendors are heading in the direction set by library--particularly Library of Congress--leadership.

Meanwhile, we need to think about how to distribute the costs, how to handle problems of authorizing system use in increasingly open systems, and how to deal with the legalities of sharing in a copyrighted world. And we need to determine what uses of the link capability will benefit our users the most and promote their development with our bibliographic networks and with the vendors who serve libraries. We are just beginning to envision the reality of linked systems and it is up to us to enlarge that vision in the service of our users.

SELECTED READINGS[1]

Avram, Henriette D. "Avram on LSP," *Library Journal* 112, no. 5 (March 15, 1987): 8, 10.

_____, and Beacher Wiggins. "The Role of the Linked Systems Project in Cooperation and Resource Sharing Among Libraries," *Libraries & Computing Centers 2* (May 1987): 1-4. (Insert of *Journal of Academic Librarianship* 13, no. 2).

Council on Library Resources, Bibliographic Services Development Program. *Technical Report on Linking the Bibliographic Utilities; Benefits and Costs*, by Donald A. Smalley [and others] Columbus, Ohio: Battelle Columbus Laboratories, 1980.

Denenberg, Ray, "NISO Draft Standard Z39.50: Information Retrieval Protocol." *Library Hi Tech News 39* (June 1987): 1, 7-8.

_____. "Open Systems Interconnection." *Library Hi Tech 3*, no. 1, Consecutive Issue 9 (1985): 15-26.

Hillkirk, John. "High-Tech World Sets Product Standards." *USA Today* (May 5, 1988): 1B.

Library of Congress Network Advisory Group. *Toward a National Library and Information Service Network: The Library Bibliographic Component.* Preliminary Edition. Washington: Library of Congress, June 1977.

The Linked Systems Project: A Networking Tool for Libraries. Edited by Judith Fenly and Beacher Wiggins. Dublin, Ohio: OCLC Online Computer Library Center, 1988.

---

[1] This list of selected readings is based on "Selected Linked Systems Project Bibliography," compiled by Beacher Wiggins of the Library of Congress and published in *The Linked Systems project: A Networking Tool of Libraries.*

McCallum, Sally H. "The Linked Systems Project: Implications for Library Automation and Networking." In *Advances in Library Automation Networking.* Edited by Joe A. Hewitt. Greenwich, Conn.: JAI Press Inc., 1987.

McGill, Michael J., Larry L. Learn, and Thomas K. G. Lydon. "A Technical Evaluation of the Linked Systems Project Protocols in the Name Authority Distribution Application." *Information Technology and Libraries*, 6, no. 4 (December 1987): 253-65.

NCLIS/NBS Task Force on Computer Network Protocol. *A Computer Protocol for Library and Information Science Applications* Washington, D.C.: National Commission on Libraries and Information Science, Dec. 1977.

Rice, James. "Open Systems: A Second Opinion." *American Libraries* 18, no. 6 (June 1987): 450-53.

ASSOCIATION OF RESEARCH LIBRARIES
1527 NEW HAMPSHIRE AVENUE N.W.
WASHINGTON, D.C. 20036
(202) 232-2466

ISBN 0-918006-14-7